

*RESPONSE RATE VIEWED AS  
ENGAGEMENT BOUTS: EFFECTS OF  
RELATIVE REINFORCEMENT AND SCHEDULE TYPE*

RICHARD L. SHULL, SCOTT T. GAYNOR, AND JULIE A. GRIMES

UNIVERSITY OF NORTH CAROLINA AT GREENSBORO

The rate of a reinforced response is conceptualized as a composite of engagement bouts (visits) and responding during visits. Part 1 of this paper describes a method for estimating the rate of visit initiations and the average number of responses per visit from log survivor plots: the proportion of interresponse times (IRTs) longer than some elapsed time (log scale) plotted as a function of elapsed time. In Part 2 the method is applied to IRT distributions from rats that obtained food pellets by nose poking a lighted key under various multiple schedules of reinforcement. As expected, total response rate increased as a function of (a) increasing the rate of reinforcement (i.e., variable-interval [VI] 4 min vs. VI 1 min), (b) increasing the amount of the reinforcer (one food pellet vs. four pellets), (c) increasing the percentage of reinforcers that were contingent on nose poking (25% vs. 100%), and (d) requiring additional responses after the end of the VI schedule (i.e., adding a tandem variable-ratio [VR] 9 requirement). The first three of these variables (relative reinforcement) increased the visit-initiation rate. The tandem VR, in contrast, increased the number of responses per visit. Thus, variables that have similar effects on total response rate can be differentiated based on their effects on the components of response rate.

*Key words:* response rate, visits, bouts, relative reinforcement, tandem variable ratio, key poke, rats

The rate of a reinforced response usually is calculated by dividing the total number of responses by the time available for the response. This method makes most sense if all instances of the response are functionally equivalent. But there are grounds for thinking that they might not be, at least not under some widely studied conditions. According to one tradition, reinforced responding is better thought of as composed of periods of engagement in the reinforced activity (i.e., visits) alternating with periods of disengagement (Blough, 1963; Gilbert, 1958; Mechner, 1992; Mellgren & Elsmore, 1991; Nevin, 1992; Nevin & Baum, 1980; Pear & Rector, 1979; Premack, 1965; Shull, 1991).

Although conventional recording equipment typically cannot distinguish visit initiations from responses that occur during the visit (but see Mechner, 1992; Nevin, 1992; Pear & Rector, 1979), these two kinds of re-

sponse might nevertheless be controlled by different variables. If so, total response rate would be a composite measure of performance rather than a unitary one. Similar changes in total response rate could result from very different patterns of change in the components. For example, one type of variable might affect total response rate by altering the visit-initiation rate, another by changing the number of responses per visit. An interesting possibility is that independent variables might be grouped on the basis of which component they primarily affect (Mechner, 1992).

To illustrate, imagine that a food-deprived rat has been obtaining food pellets intermittently by nose poking a lighted key (e.g., under a variable-interval [VI] 2-min schedule of reinforcement). Assume that the rate of poking has stabilized at about 20 per minute. To increase the rat's response rate, one could (a) increase food deprivation, (b) increase the rate of the reinforcer by decreasing the VI schedule, (c) increase the amount or taste quality of the reinforcer, (d) reduce the availability of alternative reinforcers, or (e) add a small fixed-ratio (FR) or variable-ratio (VR) schedule requirement at the end of the VI schedule (i.e., change the VI to a tandem VI FR or to a tandem VI VR). All five of these changes should increase average response

---

This research was supported by a grant from the National Science Foundation (IBN-9511934). Portions of these data were presented at the annual meetings of the American Psychological Association (1996) and the Association for Behavior Analysis (1999). Scott Gaynor is now at Western Michigan University.

Correspondence concerning this paper should be directed to R. L. Shull, Department of Psychology, Box 26164, University of North Carolina-Greensboro, Greensboro, North Carolina 27402-6164 (E-mail: rlshull@uncg.edu).

## Rat A4 - Event record of nosepokes




Fig. 1. A 30-s segment of an event record generated by a rat nose poking a lighted key. It shows the grouping of responses in engagement bouts (visits). The bracketed segment marked "a" is intended to illustrate ambiguities in determining the number of visits and their durations. A VI 1-min schedule of reinforcers had been in effect for many prior training sessions.

rate (Ferster & Skinner, 1957; Lattal, 1991). But would they do so in the same way? Perhaps not. A commonsense interpretation would group the first four variables together as motivational or incentive variables, variables that alter the propensity to engage in the reinforced activity by altering the relative reinforcement of the designated response (e.g., Herrnstein, 1970). The fifth type of variable (the added tandem ratio) does not have such an obvious commonsense interpretation of its effect. But it could be regarded as affecting mainly what the rat learns to do to obtain the reinforcer (i.e., the form of the behavioral unit) rather than the rat's propensity to initiate that unit. In this case, we might imagine that the rat learns to respond in a burst corresponding to the response requirement imposed by the tandem FR or VR schedule.

The classification of variables being suggested here is reminiscent of Morse's (1966) important distinction between the *strengthening* and the *shaping* effects of reinforcement (see also Galbicka & Platt, 1986; Logan, 1960). Our further suggestion is that these two kinds of effects might, under some conditions, be distinguished by which component of response rate changes the most. Specifically, the kinds of variables that we call *motivational* or *incentive* operations might affect the propensity to initiate visits, whereas schedule-type variables, such as adding a small VR schedule in tandem to a VI schedule, might affect the responses per visit or the response rate during a visit.

To evaluate these suggestions, it is necessary to measure the rate of visit initiations and the number of responses per visit. But that has proven to be quite difficult. The problem is that all we have to work with, ordinarily, is a string of brief responses (e.g., key pokes or lever presses) that are recorded as indistinguishable electrical impulses—as identical blips on an event record. Figure 1 shows a 30-s sample of such an event record generated by a rat that had been trained under a VI 1-min schedule of food reinforcement for nose poking a lighted key. The responses appear to occur in bouts. Indeed, this is the kind of pattern that has encouraged the two-mode view. But the record reveals some of the ambiguities associated with measuring visits. How many visits are contained in the cluster that occurs around 20 s? A reasonable case could be made that there are 1, 2, 3, 4, or even more visits, depending on how one interprets the intervals between responses. For example, is the slightly longer interresponse time (IRT) after the third response in the group best interpreted as a longer-than-usual within-visit pause, perhaps due to an unrecorded nose poke (see Palya, 1992)? Or should it be interpreted as an unusually short between-visit pause?

One approach has been to select some reasonable cutoff IRT duration (e.g., 1 s) and treat each response that ends an IRT shorter than that as a within-visit response. Then each response that ends an IRT longer than the cutoff duration is treated as a visit initiation. A problem with this approach, however, is

that the two kinds of IRT likely vary in duration, and their frequency distributions likely overlap to some extent. Thus, regardless of the cutoff duration, some unknown proportion of the responses will be misclassified. That is, some of the short IRTs that are classed as within-visit pauses might better be regarded (based, say, on functional criteria) as short between-visit pauses. And the reverse misclassification will occur with long IRTs. In recognition of this problem, some investigators report analyses using several different cutoff values in the hope of showing that the main conclusions do not depend critically on using a particular cutoff value (e.g., Mellgren & Elsmore, 1991). In addition, techniques are available for determining cutoff durations that minimize misclassifications, given certain assumptions about the frequency distributions of IRTs (Berday, 1993; Sibley, Nott, & Fletcher, 1990; Tolkamp & Kyriazakis, 1999). Such methods can increase our confidence in conclusions based on classifying visit initiations and within-visit responses based on their IRT duration.

Nonetheless, there are reasons to seek an alternative to the IRT cutoff approach. For one thing, if visits are emitted (as operant behavior generally is thought to be emitted), the frequency distribution of intervisit intervals could approximate a negative exponential distribution. Such a distribution is continuously decreasing so that its modal value is always at the shortest interval regardless of the emission rate. The emission rate is reflected in the steepness of the function's decline and in the mean IRT, but not in the location of the mode. For such a distribution, then, any particular cutoff IRT duration greater than zero will necessarily misclassify a substantial portion of the visit-initiation responses as within-visit responses. The result would be to underestimate the number of visits and to overestimate the responses per visit. Moreover, under this method the number of visits in a sample of responding and the number of responses per visit will necessarily be negatively correlated as the cutoff is varied. Thus, the two kinds of response are not assessed independently.

For the present study, we consider a different approach—one that involves no attempt to classify individual responses as visit-initiation versus within-visit responses. Instead,




Fig. 2. A diagram of the two-mode conception of response rate that was the basis for the computer simulations that generated sample IRT distributions.  $p(V)$  is the probability of initiating a visit (i.e., an engagement bout) at the end of each small time unit while disengaged;  $p(D)$  is the probability that a response during a visit will end the visit;  $1 - p(D)$  is the probability of continuing the visit following each within-visit response; and  $p(R)$  is the probability of a response per small time unit during a visit.  $p(V)$  determines the visit-initiation rate, and  $p(D)$  determines the average number of responses per visit (i.e., visit length). For the simulations,  $p(V)$  and  $p(D)$  were varied separately.  $p(R)$  was held constant at a relatively high value, thus generating a relatively high within-visit response rate.

based on characteristics of the frequency distribution of IRTs, estimates are derived for the average visit-initiation rate and for the average number of responses per visit. These two estimates are potentially independent in the sense that changes in one do not require changes in the other. Part 1 describes this technique and applies it to IRT distributions generated by computer simulations designed to show the effects of varying either visit-initiation rate or average number of responses per visit. Part 2 applies this method to IRT data from rats. The analysis shows how variables that affect total response rate also affect visit-initiation rate and average number of responses per visit. Of particular interest is whether variables that influence total response rate can be further differentiated based on which component they primarily affect.

## PART 1: COMPUTER SIMULATIONS OF RESPONDING

A simple version of the two-mode conception is summarized in Figure 2 and served as the basis for the following computer simulations of responding. For these simulations, visits were emitted randomly in time at a particular rate, responses during a visit were

emitted randomly at a much higher rate, and there was a constant probability of ending a visit immediately following each within-visit response. Two variables were manipulated. For some comparisons (identified as Sim 1), only the rate of visit initiations was varied [i.e.,  $p(V)$  at the end of each consecutive small ( $<0.15$  s) time unit during disengagement in Figure 2]. For other comparisons (identified as Sim 2), only the probability of ending a visit (and thus the average number of responses per visit) was varied [i.e.,  $p(D)$  associated with each within-visit response in Figure 2].

The purpose of the simulations was to see what the IRT distributions would look like if an experimental subject (e.g., a rat) were to respond in accordance with the simplifying assumptions expressed in Figure 2. We attempted to make the conditions of the simulations similar in a number of respects to those arranged for the rats (as will be described in Part 2). For example, the simulations were carried out in real time, using the computers and software that arranged the experimental conditions for the rats. Specifically, two components alternated every 50 s for a total of 40 cycles. One simulation operated during one of the components, and a different simulation operated during the other component. For example, the simulations might differ in the value of  $p(V)$  (Sim 1), or they might differ in the value of  $p(D)$  (Sim 2). The rate of responding during a visit [ $p(R)$  in Figure 2] was the same for all simulations; its value was set so as to generate a median within-visit IRT of 0.22 s, which was within the range of values obtained from our rats. Through some trial-and-error exploration, values of  $p(V)$  and  $p(D)$  were found that generated total response rates that were within the range of those generated by our rats.


Thus, our simulations generated samples of IRTs that were similar in size to the samples generated by our rats. Also, the simulations were constrained, as was the responding by our rats, by having 50-s components and sessions limited to 40 cycles. The distributions of IRTs from the simulations will deviate from theoretical values because of these constraints and sampling variability. As such, the simulation data should provide a useful reference for evaluating the IRT distributions produced by the rats.

There are many different ways to display frequency distributions of response times. Log survivor plots (Figure 3) are especially helpful for revealing the two-mode character of responding and for estimating visit rates and the average numbers of responses per visit. For such plots, the proportion of IRTs in the sample that are longer than some duration,  $t_i$ , is plotted as a function of  $t$  (Luce, 1986; McGill, 1963; Shull, 1991).


Because the y axis is scaled logarithmically, the slope between any two points along the x axis indicates the *relative* decrease in the frequency of IRTs between those two points. Thus, the slope indicates the frequency of responding (relative to the opportunities; cf. Anger, 1956) during the time interval bounded by the two points: The steeper the slope, the higher the relative frequency of responding during the interval. If the frequency of responding relative to opportunities is the same regardless of elapsed time since a response, then the log survivor plot will be a decreasing straight line, approximating a simple exponential decay. By the same reasoning, if the relative frequency of responding increases as a function of elapsed time, the log survivor function will curve downward. And if the relative frequency of responding decreases as a function of elapsed time, the function will be decreasing, concave upward.

In short, to the extent that responding approximates a random emitter, the log survivor data will plot as a decreasing straight line. The steeper the line, the higher the emission rate is.

The two log survivor plots in the top panel of Figure 3 come from simulations that differed only in the rate of initiating visits (Sim 1). The striking feature of these plots is their broken-stick appearance (Clifton, 1987). Such a plot suggests that two different response-generating modes are operating (Slatyer & Lester, 1982; Tolkamp & Kyriazakis, 1999). Indeed, we know this to be the case for the plots in Figure 3 (top) because the simulations were programmed based on the two-mode conception. The initial steeply declining segment in each plot represents within-visit responding, which occurred at a high rate. The more gradually declining limb in each plot represents visit initiations. The slopes of the visit-initiation limbs differ between the two plots, which was expected be-



	Total Rs/m	Visits/m	Rs/visit
<u>A</u> (Lo visit rate)	19.7	8.9	2.9
<u>B</u> (Hi visit rate)	52.4	22.2	2.8



	Total Rs/m	Visits/m	Rs/visit
<u>A</u> (Short visits)	19.7	8.9	2.9
<u>B</u> (Long visits)	50.2	8.8	8.0

Fig. 3. Log survivor plots from computer simulations. These plots show the proportion of IRTs (logarithmic scale) longer than some duration,  $t$ , as a function of elapsed time,  $t$ , since a response. The top panel presents

cause the two simulations differed in visit-initiation rates. Moreover, the limbs are reasonably linear (on the logarithmic  $y$  axis), which was expected because the probability of a visit initiation for each simulation was constant as a function of elapsed time during disengagement (but differed between the pair of simulations).

If the visit-initiation limbs on the right (Figure 3) are extrapolated back to the point at which they would intersect the  $y$  axis, that point indicates the proportion of all responses that are visit initiations. The two plots in the upper panel appear to intersect the  $y$  axis at about the same point. This, too, is as it should be because the probability of ending a visit—and, thus, the average number of responses per visit—was set to be the same for these two simulations.

The two log survivor plots in the lower panel of Figure 3 (Sim 2) show a different comparison. For these two simulations, the visit-initiation rate was set the same but the probability of ending a visit differed. Consequently, the visit-initiation limbs are approximately parallel, indicating approximately equivalent visit-initiation rates. But the points at which these limbs would intersect the  $y$  axis differed. The lower the intersection point is, the lower the proportion of all responses that are visit initiations and thus the higher the average number of responses per visit.

Thus, the two sets of log survivor plots show the effect of changing only the visit rate (upper panel) or changing only the proba-

the results from Sim 1, in which the two simulations differed only in the probability of initiating a visit (i.e., in the visit-initiation rate). The lower panel presents the results from Sim 2, in which the two simulations differed only in the probability of terminating a visit (i.e., in the average number of responses per visit). (Note that Plot A is the same in the upper and lower sets.) To estimate visit-initiation rate and average number of responses per visit, straight lines were fit (Pearson least squares) to the natural logarithms of the segment of the plots ranging from  $t = 1$  s through 12 s (shown by the straight line through the filled and open circles). The equation from each fit is shown in its exponential form. The numbers in the chart under each set of plots indicate for each plot the total number of responses per minute calculated the usual way from response totals, the visit-initiation rate in visits per minute (the numerical value of the exponent  $\times 60$ ), and the average number of responses per visit ( $1/\text{coefficient}$ ).



bility of ending the visit (lower panel). That there are two modes of responding is apparent from the sharp break in each plot. Moreover, the slope of the limb to the right of the break in each plot and the  $y$ -axis intercept of the limb, extrapolating the limb back to the  $y$  axis, have convenient interpretations. The slope of the limb indicates the visit-initiation rate, and the  $y$ -axis intercept indicates the proportion of responses that are visit initiations, the inverse of which is the average number of responses per visit.

When the two modes of responding occur at very different rates, and when the limbs are reasonably linear on a logarithmic  $y$  axis, there is a simple method for generating numerical estimates of visit-initiation rate and responses per visit: Fit a straight line to a stable segment of the visit limb of a log survivor plot. Because the  $y$  axis is scaled logarithmically, the fits should be to the logarithms of the  $y$ -axis values of the limb segment. We have used the segment from 1 s through 12 s for our fits. Although somewhat arbitrary, this range seemed to work well for our data sets. The 1-s lower boundary was sufficiently to the right of the break to contain visit-initiation responses. And beyond 12 s, the number of IRTs in the sample was often very small, thus producing unstable values. The fits to these limb segments give the slope and  $y$ -axis intercept.

Figure 3 provides examples of such fits. Also shown in each panel are the best fitting equations, expressed in exponential form [i.e.,  $y = a \cdot \exp(b \cdot t)$ ]. Using the exponential form, the value of  $a$ , the  $y$ -axis intercept, can be read directly as the proportion of responses that are visit initiations. In addition, because the  $x$  axis is scaled in terms of seconds, the number in the exponent (disregarding sign) indicates the visit initiations per second.<sup>1</sup>

<sup>1</sup> Expressed in logarithmic form [i.e.,  $\ln(y) = \ln(a) + b \cdot t$ ] the value,  $\ln(a)$ , would have to be converted to the antilog to determine the proportion. One further note is that the fits are best made by converting the original survivor points to natural logarithms (base  $e$ ) rather than common logarithms (base 10). The reason is that in exponential form the  $b$  term in the exponent can be read directly (ignoring the negative sign) as a visit-initiation rate, in units of  $t$ . If the fits are to the common logarithms, an additional (albeit simple) transformation is required to read the value of  $b$  as a response rate in units of  $t$ . (We use the exponential trendline routine in Micro-

The chart below each panel in Figure 3 summarizes the measures of responding for each plot in the pair. By design, for each pair the difference in total response rate between simulations is about the same. This difference, however, is generated through different patterns of change in the components of response rate—namely, through differences in visit-initiation rate for the pair in the top panel and through differences in the average number of responses per visit for the pair in the lower panel.

Because the simulation program tagged (i.e., coded) the responses as either visit initiations or within-visit responses, it was possible to calculate visit-initiation rates and the average number of responses per visit directly from these coded response and time totals for the session. Visit-initiation rate was calculated by dividing the number of visit initiations (determined from the coded responses) by the time available for visit initiations. That time was the session time (for the relevant component) minus the time spent engaged in visits. The average number of responses per visit was calculated by dividing the total number of responses by the number of visit initiations.

The scatter plots in Figure 4 indicate that the estimates derived from the log survivor fits from 10 simulations corresponded fairly well to values derived from direct calculations based on the coded responses. This correspondence indicates that the fits to the limbs of log survivor plots can provide valid estimates of visit-initiation rate and average number of responses per visit. It is apparent, however, that the visit-initiation rates were consistently higher when estimated from the log survivor fits than when calculated directly from the coded response data (i.e., the solid points in the top panel are above the diagonal that indicates exact correspondence). In part, this discrepancy results from the fact that the slope of the log survivor limb estimates the average visit-initiation rate that would occur if the whole time were available

soft Excel® for these fits. The logarithmic transformations [base  $e$ ] are carried out automatically.) Note that the estimates will likely be quite different if exponential functions are fitted to the unlogged proportions. We think it makes most sense to fit a straight line to the logarithms because doing so corresponds most closely to how the log survivor plots appear.

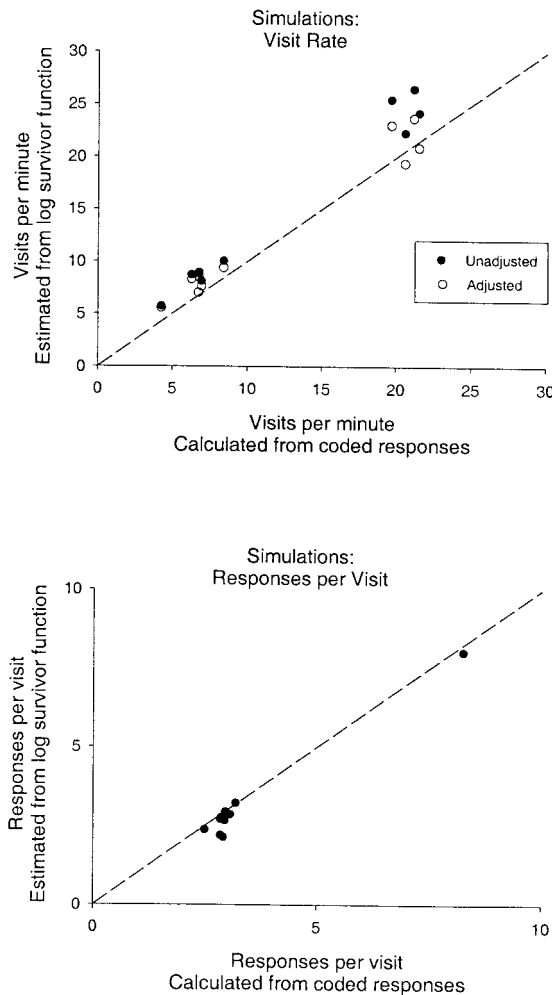


Fig. 4. Scatter plots, based on 10 simulations, showing the correspondence between two ways of computing visit-initiation rate (top panel) and average number of responses per visit (bottom panel). The performance measures estimated from the tails of the log survivor plots (as illustrated in Figure 3) are plotted as a function of the corresponding performance measures computed directly from coded responses. The diagonal lines indicate perfect correspondence. The filled points in the top panel indicate the estimated visit-initiation rates based on the numerical values of the exponents (as in Figure 3). The open points indicate the estimated visit-initiation rate after adjusting the values for proportion of time occupied by the visits, as described in the text.

for visit initiations. But a visit cannot be initiated while a visit is already engaged. Thus, to produce a closer correspondence between the estimated visit rate (from the log survivor limbs) and the observed visit rate (here, from the coded responses), it is necessary to adjust the estimated visit rate downward, taking into

account visit-engagement time. As discussed more fully in the Appendix, the adjusted visit-initiation rate,  $V'$ , was computed by

$$V' = V \left( \frac{(1/V)}{(1/V) + NW} \right), \quad (1)$$

where  $V$  indicates the unadjusted visit-initiation rate,  $N$  indicates the average number of responses per visit, and  $W$  indicates the within-visit average IRT. We set  $W$  equal to the median within-visit IRT for the simulations (i.e., 0.0037 min or 0.22 s). The adjusted visit-initiation rates (shown as the open points in the top panel of Figure 4) correspond closely to the observed values calculated from the visit-initiation totals as coded by the simulations (albeit with a small remaining upward displacement from the line of exact correspondence that we do not yet understand).

We think the unconstrained estimates of visit-initiation rates are preferable to the actually measured visit-initiation rates because the latter are constrained, and thus limited in generality by the duration of the response unit, namely visits. For that reason, in the remainder of the paper we will report estimates of visit-initiation rates that are not adjusted. In any case, the differences between the unconstrained and adjusted values will usually be small except when visit duration is long and visit-initiation rate is high (see Shull, 1991, pp. 263–265).

The samples of IRTs were of sufficient size to generate valid and stable estimates of visit-initiation rate and average number of responses per visit. It should be noted, however, that log survivor plots with small sample sizes will tend to be variable and will lead to inconsistent estimates. Our daily samples (from the simulations and from the rats) typically contained between 400 and 2,000 responses; some samples, however, contained as few as 200 responses, and for Rat A3 (below) there were occasionally as few as 50.

## PART 2: DATA FROM RATS

Part 1 described a method for computing two different dependent variables from fits to segments of log survivor plots of IRT distri-

butions. The data from the simulations showed that, under some conditions at least, the slope and intercept of the limb segment could be interpreted as valid estimates of visit-initiation rate and average number of responses per visit, respectively. An important question at this point is whether these new dependent variables vary systematically as a function of independent variables so as to provide new insights about the control of behavior.

## METHOD

### *Subjects*

The subjects were 7 male Long Evans hooded rats, identified as members of Squad A (4 rats) or Squad B (3 rats; a 4th rat in this squad died before the start of the present project). They were obtained from a commercial breeder at about 6 weeks of age (about 150 g) and were housed in individual plastic cages throughout the project. Water was freely available in each rat's home cage and in the experimental chamber. The housing rooms were maintained at about 22 °C and were on a 12:12 hr light/dark cycle; the experimental sessions were conducted during the light periods.

For the first several weeks, the rats were handled daily to adapt them to being picked up and carried. Following this acclimation period, a food-deprivation regimen was begun. Each rat was allowed free access to food (lab blocks) for about 1.5 hr each day. Training in the experimental chambers began when the rats had grown to about 250 g (about 4 months old). Once training began, the rats were given free access to food in their home cages for at least 1 hr shortly after each daily session (usually, between 1 and 1.5 hr). With this regimen, the rats continued to gain weight, gradually leveling off at between 315 and 345 g—a level that was maintained throughout the remainder of the project. This weight level is the same as described by Ator (1991) as an appropriate and effective deprivation level for male Long Evans hooded rats.

The rats differed in age and experimental histories. At the beginning of this project, the rats in Squad A were about 15 months old and those in Squad B were about 7 months old. The rats in Squad A had spent the pre-

vious year in a study involving choice (by pressing one of two levers) between different signaled delays of food reinforcement; the rats in Squad B had spent the previous 4 months in a study involving key poking reinforced by food pellets under various multiple schedules of food reinforcement.

### *Apparatus*

The experimental chambers were four similar two-lever operant chambers (30 cm wide by 32 cm deep by 30 cm high) constructed of sheet metal (top and three sides) and clear plastic (rear door, 21 cm by 30 cm). The floor was made of stainless-steel rods (0.7 cm diameter) spaced 1 cm apart. In the middle of the front panel, centered 4.3 cm above the floor, was a rectangular opening (4.4 cm by 4.4 cm) giving access to a small metal food tray. An electrical pulse to a Gerbrands pellet dispenser located behind the front panel caused a food pellet (45 mg Noyes Formula A) to drop into the food tray. The operation of the dispenser made a click, and the pellet landing in the tray made a plinking sound a fraction of a second later (both sounds could be easily heard by a person standing several meters from the chamber). Reinforcer deliveries were not signaled by any change in illumination. Each rat always had its daily session in the same chamber.

A translucent plastic key (Lehigh Valley Electronics) was mounted behind a round hole (1.9 cm diameter) through the left wall. The center of the key was 5.1 cm toward the rear of the chamber from the front wall and 6.2 cm above the floor. A recordable response occurred when the key was pushed a distance of about 1.5 mm (measured at the center) with a force of at least 0.18 N (three of the four chambers) or 0.3 N (the fourth chamber; Rats A4 and B3). Such responses produced a brief click from a small snap-action switch connected to the key. The rats were nearly always observed to press the key with their noses; occasionally they used their paws. The key could be illuminated from behind with either a blinking white light (two flashes per second) or a continuously illuminated white light (2 GE 1829 bulbs operated at 28 V DC).

Toward the back of the chamber along the left wall (approximately 24 cm from the front wall and 2.5 cm above the floor) was a small



hole through which a metal drinking spout protruded about 1 cm into the chamber. The spout was attached to a water bottle suspended outside the chamber and allowed free access to water throughout each session.

The four chambers were placed on shelves on a cart, two on the top shelf and two on a lower shelf. No attempt was made to shield the chambers from sounds from other chambers. It was apparent that each rat quickly learned to go to its food tray only when its own feeder operated. Normally, the four chambers operated at the same time (although out of phase). But occasionally, as a check, a daily session was conducted with one of the rats in isolation. No systematic difference was detected as a function of whether or not the other chambers were operating.

During experimental sessions the lights in the room were darkened, but we could observe the rats easily when the keys were lighted and even when the keys were darkened because of dim light that entered the room from the hallway through a translucent window in the door. There was no other source of illumination in the chambers (i.e., no houselight or feeder light).

Each chamber was connected to its own dedicated special-purpose computer (Walter & Palya, 1984) for controlling the experimental events and recording the data.

#### *Procedure*

For all 10 conditions the basic procedure was a multiple schedule. That is, during each session the keylight alternated every 50 s between blinking (two flashes per second) and steady illumination (with a 5-s dark period following each keylight presentation). Each daily session consisted of 40 blink-steady cycles. Throughout a block of daily sessions (i.e., within a condition) one schedule of reinforcement was correlated with the blinking keylight and a different schedule of reinforcement was correlated with the steady keylight. Responding had no programmed consequences during the 5-s blackouts between keylight components.

Our intent was to generate two different performances per condition, one under discriminative control of the blinking light and the other under discriminative control of the steady light. One of the discriminative stimuli in each pair (the rich-reinforcement signal)

was correlated with about four times as many pellets per cycle as the other discriminative stimulus (the lean-reinforcement signal). But within this constraint, the schedules were altered in various ways between conditions in the hope of producing a variety of different response rates.<sup>2</sup> (Table 1 summarizes these conditions; detailed descriptions of the schedules are presented below.) The IRTs generated by these various schedules were then used to assess the effects of the different independent variables on visit-initiation rate and average number of responses per visit.

The set of 10 conditions permitted comparison of two levels of the following independent variables: (a) rate of reinforcement (15 reinforcers per hour vs. 60 reinforcers per hour), (b) amount of reinforcement (one food pellet per reinforcer vs. four food pellets), (c) percentage of reinforcers that were contingent on key poking (25% vs. 100%), and (d) the imposition of a tandem response requirement at the end of a VI schedule (no additional requirement vs. a tandem FR or VR requirement). Some of these comparisons were made between the two discriminative stimuli of a multiple-schedule pair; other comparisons were made between conditions (see Table 2).

Interresponse times were recorded separately for each discriminative stimulus during the last three sessions of each condition. The computer recorded IRTs in units of 0.01 s, but we suspect that our system actually resolved IRTs to only about the nearest 0.1 s. The IRT for the first response in each 50-s discriminative-stimulus component was measured from the start of the component; the IRTs of all subsequent responses during the component were measured from the previous response.

A more detailed description of each of the schedule pairs follows. The conditions are identified by their order in the series of 10 conditions (Table 1). Unless noted otherwise, the reinforcer was a single 45-mg food pellet.

For Condition 1, the rich-reinforcement discriminative stimulus (blinking light) was

<sup>2</sup> One or two resistance-to-extinction tests were given after each of these conditions to see if there was any relation between visit-initiation rate and resistance to extinction in the presence of the discriminative stimuli (see Nevin & Grace, 2000). We hope to present data from some of these extinction tests in a subsequent paper.

Table 1

The order of the 10 conditions, the schedules of reinforcement that were in effect in the presence of each of the two discriminative stimuli during each condition, the type of discriminative stimulus that signaled the rich-reinforcement component (blinking keylight [B] or steadily illuminated keylight [S]), and the number of sessions devoted to each. The rich-reinforcement component of the multiple-schedule pair provided an average of about four times as many food pellets per component as did the lean-reinforcement component. Replications are indicated in parentheses.

Condition order	Schedules of multiple pair		Stimulus rich	Number of sessions
	Rich	Lean		
1	VI 1 min	VI 4 min	B	27 (Squad A) 30 (Squad B)
2	VI 4 min + VT 1.33 min	VI 4 min	B	30
3	VI 2 min: 4 pellets	VI 2 min: 1 pellet	S	34
4	VI 2 min: 4 pellets	Tandem VI 2 min FR 4: 1 pellet	S	30
5	VI 2 min: 4 pellets	Tandem VI 2 min VR 9: 1 pellet	S	17
6 (3)	VI 2 min: 4 pellets	VI 2 min: 1 pellet	S	29
7	VI 2 min: 1 pellet + VT 2 min: 3 pellets	VI 2 min: 1 pellet	S	43
8	VI 2 min: 1 pellet + VT 2 min: 3 pellets	Tandem VI 2 min VR 4: 1 pellet	S	26
9	VI 2 min: 1 pellet + VT 2 min: 3 pellets	Tandem VI 2 min VR 9: 1 pellet	S	24
10 (7)	VI 2 min: 1 pellet + VT 2 min: 3 pellets	VI 2 min: 1 pellet	S	25

correlated with a VI 1-min schedule, and the lean-reinforcement discriminative stimulus (steady light) was correlated with a VI 4-min schedule. Each of the VI schedules (here and throughout the experiment) was comprised of 16 different intervals and provided a roughly constant probability of reinforce-

ment in time since the last reinforcer (Fleshler & Hoffman, 1962; Hantula, 1991). Upon the first presentation of a component in a session and immediately following each presentation of an assigned reinforcer, one of the 16 intervals was randomly selected (with replacement). Once an interval had assigned a

Table 2

The source of comparisons for determining the effects of types of independent variables on aspects of responding. The left column shows the type of variable being compared. The right column shows the condition (identified by the condition order shown in Table 1) and the component of that condition (rich reinforcement vs. lean reinforcement).

Type of comparison, expected to generate low total response rate (left) vs. high response rate (right)	Conditions compared
Rate of reinforcement (15 vs. 60 reinforcers per hour)	
VI 4 min vs. VI 1 min	1-Lean vs. 1-Rich
Amount of reinforcement (1 vs. 4 pellets)	
VI 2 min: 1 pellet vs. VI 2 min: 4 pellets	3-Lean vs. 3-Rich 6-Lean vs. 6-Rich
Percentage contingent reinforcement (25% vs. 100%)	
[VI 4 min + VT 1.33 min] vs. VI 1 min	2-Rich vs. 1-Rich
[VI 2 min: 1 pellet + VT 2 min: 3 pellets] vs. VI 2 min: 4 pellets	7-Rich vs. 3-Rich 10-Rich vs. 6-Rich
Tandem VR (1 vs. 9 responses)	
VI 2 min vs. [tandem VI 2 min VR 9]	3-Lean and 6-Lean vs. 5-Lean 7-Lean and 10-Lean vs. 9-Lean

reinforcer, the next interval did not begin to time until the assigned reinforcer had been delivered. If an assigned reinforcer had not been delivered when the current 50-s multiple-schedule component ended, the reinforcer assignment was saved and carried over to the next presentation of that component. Likewise, at the end of a 50-s component, the time elapsed during an interval of a VI schedule was saved so that the interval resumed timing at the same point when the program returned to that component.

For Condition 2 during the rich-reinforcement discriminative stimulus, key pokes were reinforced by a VI 4-min schedule (15 pellets per hour) while a VT 1.33-min schedule operated concurrently delivering pellets independently of any response (45 pellets per hour), for a total of 60 reinforcers per hour. Thus, approximately 25% of the reinforcers were contingent on key poking. The lean-reinforcement discriminative stimulus again was correlated with a VI 4-min schedule of reinforcement (i.e., 15 reinforcers per hour). The VT schedules were constructed and operated as described above for the VI schedules.

Under Condition 3, both discriminative stimuli were correlated with a VI 2-min schedule. The number of pellets presented per reinforcement differed in the two components. In the presence of the rich-reinforcement discriminative stimulus, the reinforced response produced four pellets (the first coming immediately after the response and the next three following at 0.25-s intervals). For the lean-reinforcement discriminative stimulus, the reinforced response produced a single pellet, as before. These schedules are designated as VI 2 min: 4 pellets and VI 2 min: 1 pellet.

Under Condition 4 the rich-reinforcement discriminative stimulus was again correlated with a VI 2-min: 4-pellet schedule. The lean-reinforcement discriminative stimulus was correlated with a tandem VI 2-min FR 4: 1-pellet schedule. That is, for a response to be reinforced, the time interval specified by the VI 2-min schedule had to elapse and then five responses had to occur (rather than just the one response normally required by a VI schedule). No stimulus change indicated the completion of the VI interval and the start of the FR 4 requirement. There is some debate

about why adding a small FR schedule in tandem with a VI schedule might increase response rate, but there is no doubt that it does so (Ferster & Skinner, 1957; Killeen, 1969; Peele, Casey, & Silberberg, 1984). Moreover, if the size of the tandem FR is fairly small, the rate of reinforcement will remain close to the maximum allowed by the VI schedule.

Condition 5 was identical to Condition 4 except that the additional tandem response requirement was a VR 9 schedule of reinforcement instead of an FR 4. The response requirement of the VR ranged from a minimum of 1 to a maximum of 17. Each response requirement within this range had an equal probability of occurring (i.e., the frequency distribution was rectangular).

Condition 6 was a replication of Condition 3.

Under Condition 7, in the presence of the rich-reinforcement discriminative stimulus, a VI 2-min: 1-pellet schedule was in effect for key poking, and a VT 2-min: 3-pellet schedule operated concurrently. Again, the first pellet of the three was delivered immediately after the VT interval had elapsed, and the other two were delivered after successive intervals of 0.25 s. The lean-reinforcement discriminative stimulus was correlated with a VI 2-min: 1-pellet schedule.

Conditions 8 and 9 were identical to Condition 7 except that during the lean-reinforcement discriminative stimulus a VR 4 schedule (Condition 8) or a VR 9 schedule (Condition 9) was added in tandem to the VI 2-min: 1-pellet schedule.

Condition 10 was a replication of Condition 7.

With few exceptions, sessions were conducted 7 days per week. During the first few conditions, response rates appeared to stabilize by about 15 training sessions. We therefore typically conducted at least 20 training sessions per condition (see Table 1), stopping at a convenient point thereafter provided that there was no apparent upward or downward trend over sessions in total response rate for any rat in the presence of either discriminative stimulus.

## RESULTS

One-session samples of log survivor plots from the 4 rats in Squad A (Figure 5) suggest that, indeed, the slope and intercept of the

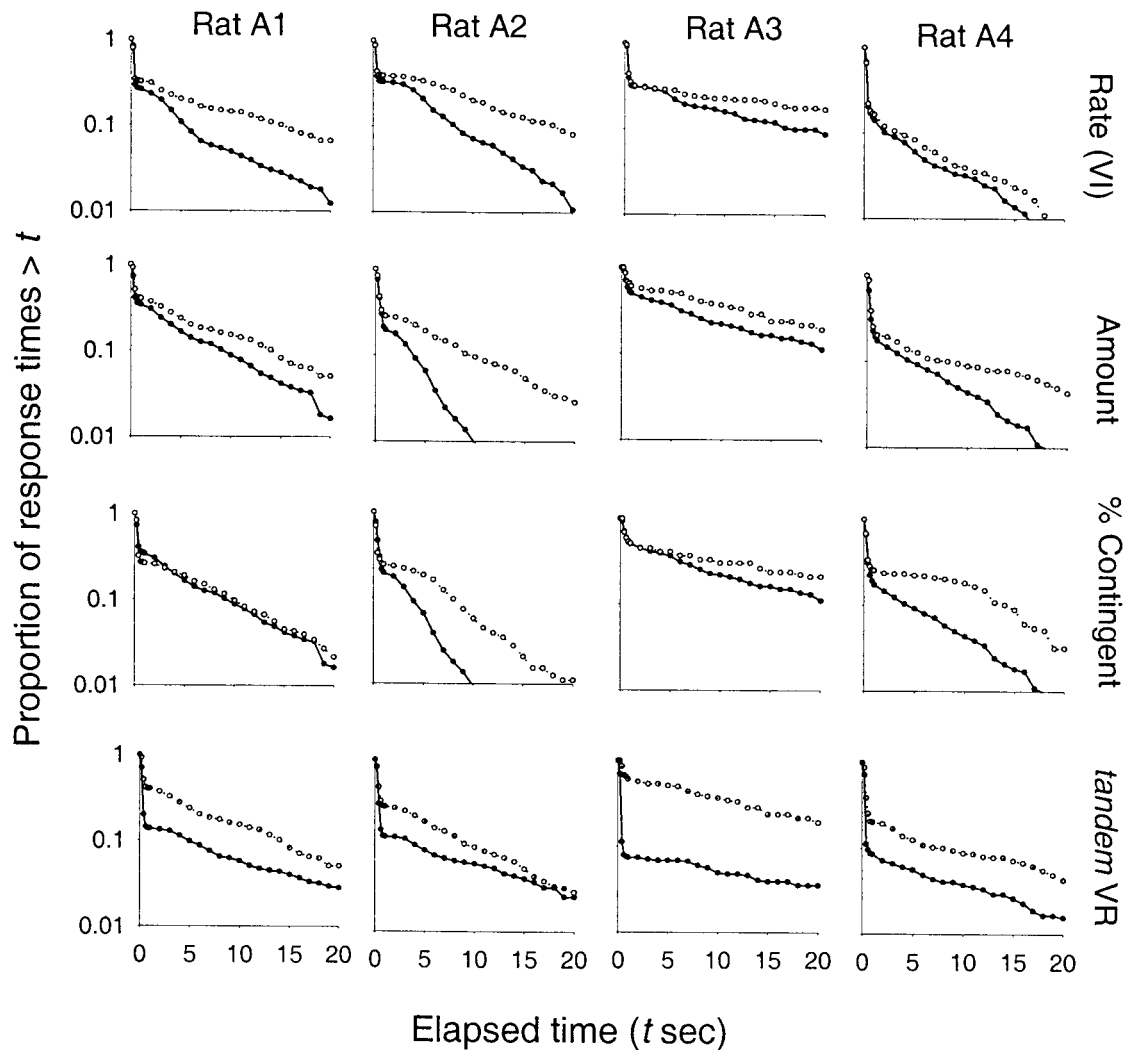


Fig. 5. Single-session log survivor plots from the last sessions of the conditions indicated for each of the rats in Squad A. Proportions are plotted for consecutive 0.2-s intervals from 0 through 1 s; they are plotted for consecutive 1-s intervals from 1 s through 20 s. The following summarizes the specific comparisons (Tables 1 and 2). Top row: VI 4 min (open) versus VI 1 min (filled) (i.e., Condition 1-Lean vs. Condition 1-Rich). Second row from the top: VI 2 min: 1 pellet (open) versus VI 2 min: 4 pellets (filled) (i.e., Condition 3-Lean vs. Condition 3-Rich). Third row from the top: 25% of the 60 reinforcers per hour response contingent (open) versus 100% response contingent (filled) (i.e., Condition 2-Rich vs. Condition 1-Rich). Bottom row: the absence (open) or presence (filled) of a tandem VR 9 requirement at the end of a VI 2-min schedule (i.e., Condition 3-Lean vs. Condition 5-Lean).

log survivor limbs are affected by different classes of independent variables. Although there were some exceptions, differences in the rate of reinforcement, amount of reinforcement, and percentage of reinforcers that were response contingent affected the slope of the right limb (i.e., what we interpret as estimating visit-initiation rate) but had little systematic effect on the y-axis intercept (see plots in the top three rows of Figure 5).

Adding a tandem VR 9 schedule (bottom row), in contrast, strongly affected the y-axis intercept of the limb (the inverse of which we interpret as estimating the average number of responses per visit). The tandem VR sometimes had an additional effect in that the visit-initiation limb was sometimes less steep with the tandem VR than without (cf., e.g., the plots for Rat A2, bottom row). The implication is that adding the tandem VR re-

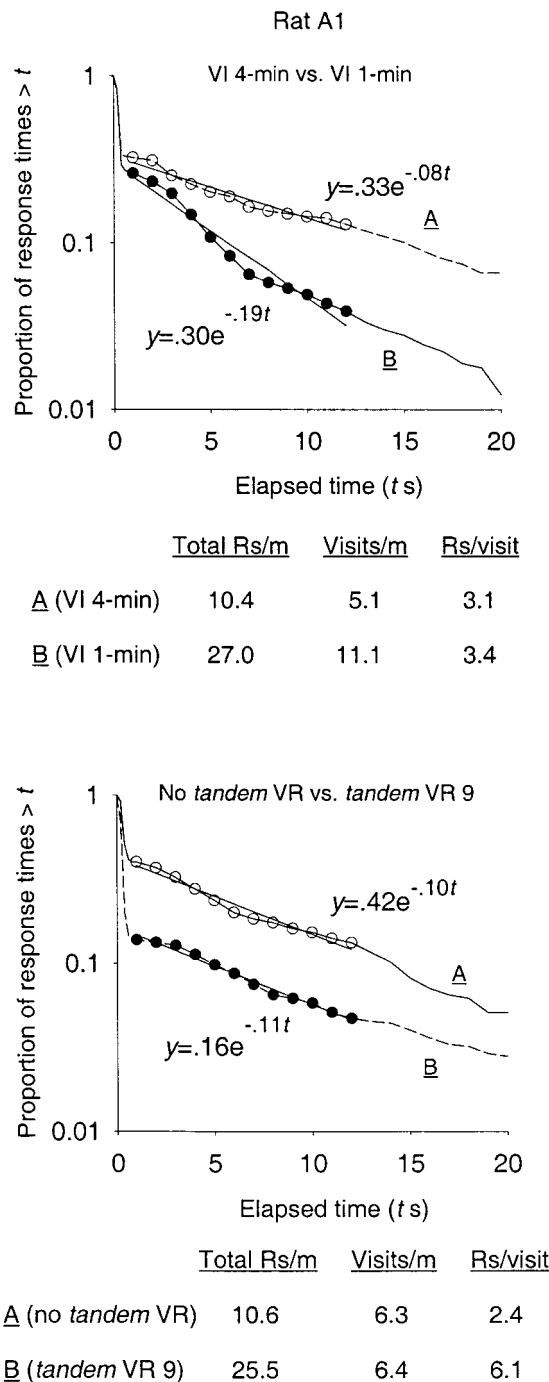


Fig. 6. Illustration of the method of estimating visit-initiation rate and average number of responses per visit from log survivor plots generated by Rat A1. Each IRT distribution (plot) comes from a single session. The pair of plots in the upper panel show the effect of varying the rate of reinforcement (VI 4 min vs. VI 1 min). The pair of plots in the lower panel show the effect of adding a VR 9 schedule in tandem to a VI 2-min schedule. As de-




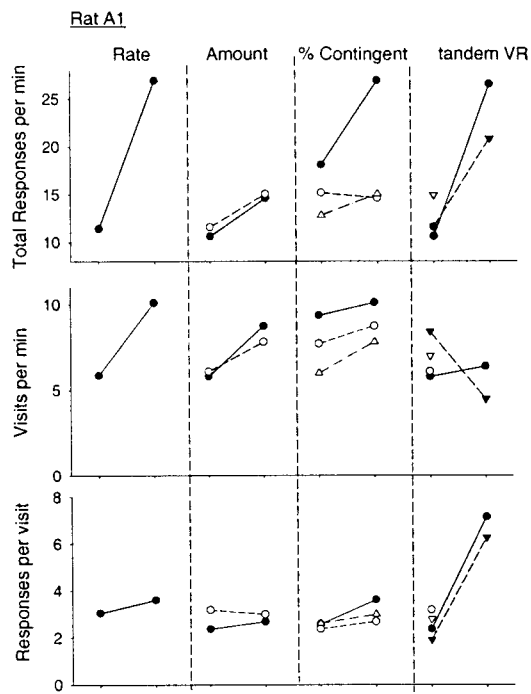
quirement increased the length of a visit (responses per visit) but reduced the rate of initiating visits.

Although visual inspection of these plots indicates that the different classes of experimental variables affected the characteristics of the log survivor plots differently, it would be helpful to have quantitative estimates of these effects. Figure 6 illustrates the application of the method for estimating visit-initiation rate and number of responses per visit (discussed earlier in relation to the simulation data; Figure 3) to data generated by 1 of the rats. The top panel shows the effects of different rates of reinforcement (VI 1 min vs. VI 4 min) on the log survivor plots. The most obvious effect was on the slope of the right limb, an effect that we interpret as changing the visit-initiation rate. The bottom panel shows that including or not including a tandem VR 9 schedule at the end of a VI 2-min schedule affects the y-axis intercept of the limb, an effect that we interpret as changing the number of responses per visit. The limbs in both sets are reasonably linear (on the logarithmic y axis) so that the fits to the limb segments (1 to 12 s) seem to characterize the limbs reasonably well. The estimated visit-initiation rates and responses per visit, listed in the chart under each set of plots, confirm that rate of reinforcement affected the visit-initiation rate and that adding the tandem VR affected the responses per visit. Both kinds of manipulation had about the same effect on total response rate (listed in the charts), but they affected the components of response rate quite differently.

This kind of analysis was applied to the log survivor plots for all the rats. The fits were carried out on the IRT data from each of the last three sessions of each condition. The median values of visit-initiation rate and responses per visit, as determined from these fits, are plotted for the indicated conditions for the rats in Squad A (Figure 7) and Squad B (Figure 8). Also shown are the corresponding to-

scribed for Figure 3, straight lines have been fitted to the natural logarithms of the segment of the plots ranging from  $t = 1$  s through 12 s (shown by the straight line through the filled and open circles) and the equations listed. The chart under each set of plots (see Figure 3) indicates the relevant performance measures.





tal response rates (medians) for the indicated conditions. As expected, total response rate (top row in each panel) was increased by increasing the rate of reinforcement (from VI 4 min to VI 1 min), the amount of reinforcement (from one to four pellets per reinforcer), the percentage of reinforcers that were response contingent (from 25% to 100%), and by adding a VR 9 component in tandem with a VI schedule. These variables, however, had different effects on the components of total response rate, consistent with the impressions from inspecting the log survivor plots. Increases in reinforcement rate, amount, and percentage of response-contingent reinforcement increased visit-initiation rate (middle row, first three columns of each panel). These variables also tended to increase the number of responses per visit (bottom row, first three columns), but this effect was smaller and less consistent than the effect on visit-initiation rate. In contrast, adding the tandem VR 9 requirement increased the number of responses per visit (bottom row, far right in each panel) and often decreased the visit-initiation rate (middle row, far right).

One notable feature of the data in Figures 7 and 8 is that the trends held up over replications, despite the fact that the schedule correlated with the alternative discriminative stimulus of the multiple-schedule pair sometimes differed between replications. The consistency across replications appeared to be about the same, regardless of the measure of performance. Moreover, the trends appeared to be reasonably consistent across the 7 rats.

Although not critical to our conclusion that the tandem VR 9 strongly affected visit length, we nevertheless wondered to what extent the long visit length (i.e., the large number of responses per visit) under the tandem VI 2-min VR 9 schedule represented a stable (i.e., learned) characteristic of performance under control of the discriminative stimulus. Figure 9 addresses this question by showing, for the rats in Squad A, the log survivor plots from the last session under the tandem VI 2-min VR 9 schedule (Condition 5-Lean) and from the very next session when the tandem VR 9 component was removed (Condition 6-Lean). Also shown, for comparison, are the plots from the last sessions under Condition 3-Lean and Condition 6-Lean (i.e., a VI 2-min schedule without the tandem VR). The important finding is that the average number of responses per visit on the first session without the tandem VR 9 component was similar to that on the last session with the tandem VR 9 component. Apparently, the long visit length that is engendered by the tandem VR component persists after the VR component is removed. In other words, the long visit length appears to be a learned response pattern rather than an artifact of whatever response requirement is currently in effect. The results were similar for the rats in Squad B (see Table 3, which shows, for all the rats, the average number of responses per visit estimated from the log survivor plots for these sessions).

Figures 7 and 8 (bottom right set for each rat) showed that adding a tandem VR 9 re-

←

Fig. 7. The effects of different independent variables on total response rate, visit-initiation rate, and responses per visit for each of the 4 rats in Squad A. Each set of panels presents the data from a different rat. Each data point is the median value of the last three sessions of the condition. Total response rate (top row in each panel) was computed from response totals. Visit-initiation rate (middle row) and the average number of responses per visit (bottom row) were estimated from the limbs of the log survivor plots. Data are shown for two levels of each variable. For each type of variable, the x-axis value on the left was expected to generate the lower total response rate, and the x-axis value on the right was expected to generate the higher total response rate. Replications are indicated by dashed lines. Unconnected points for the tandem comparison (VI alone) indicate additional replications of the VI 2-min schedule with one pellet per reinforcer. Although it is not essential to appreciate the trends, one can determine the particular conditions and components that generated the points from Table 2. In some cases, the pair of points (left vs. right) came from the two different multiple-schedule components of a single condition. Other pairs came from multiple-schedule components from two different conditions. Moreover, the percentage-of-contingent-reinforcement variable was changed two different ways (see Tables 1 and 2). For one comparison (solid line), the 25% value was generated by a VI 4-min plus VT 1.33-min schedule during the discriminative stimulus (Condition 2-Rich) and the 100% value was generated by a VI 1-min schedule (Condition 1-Rich). For the other two comparisons, the 25% value was generated by a schedule arranging VI 2 min: 1 pellet plus VT 2 min: 3 pellets (Condition 7-Rich and Condition 10-Rich), and the 100% value was generated by a VI 2-min: 4-pellet schedule (Condition 3-Rich and Condition 6-Rich).

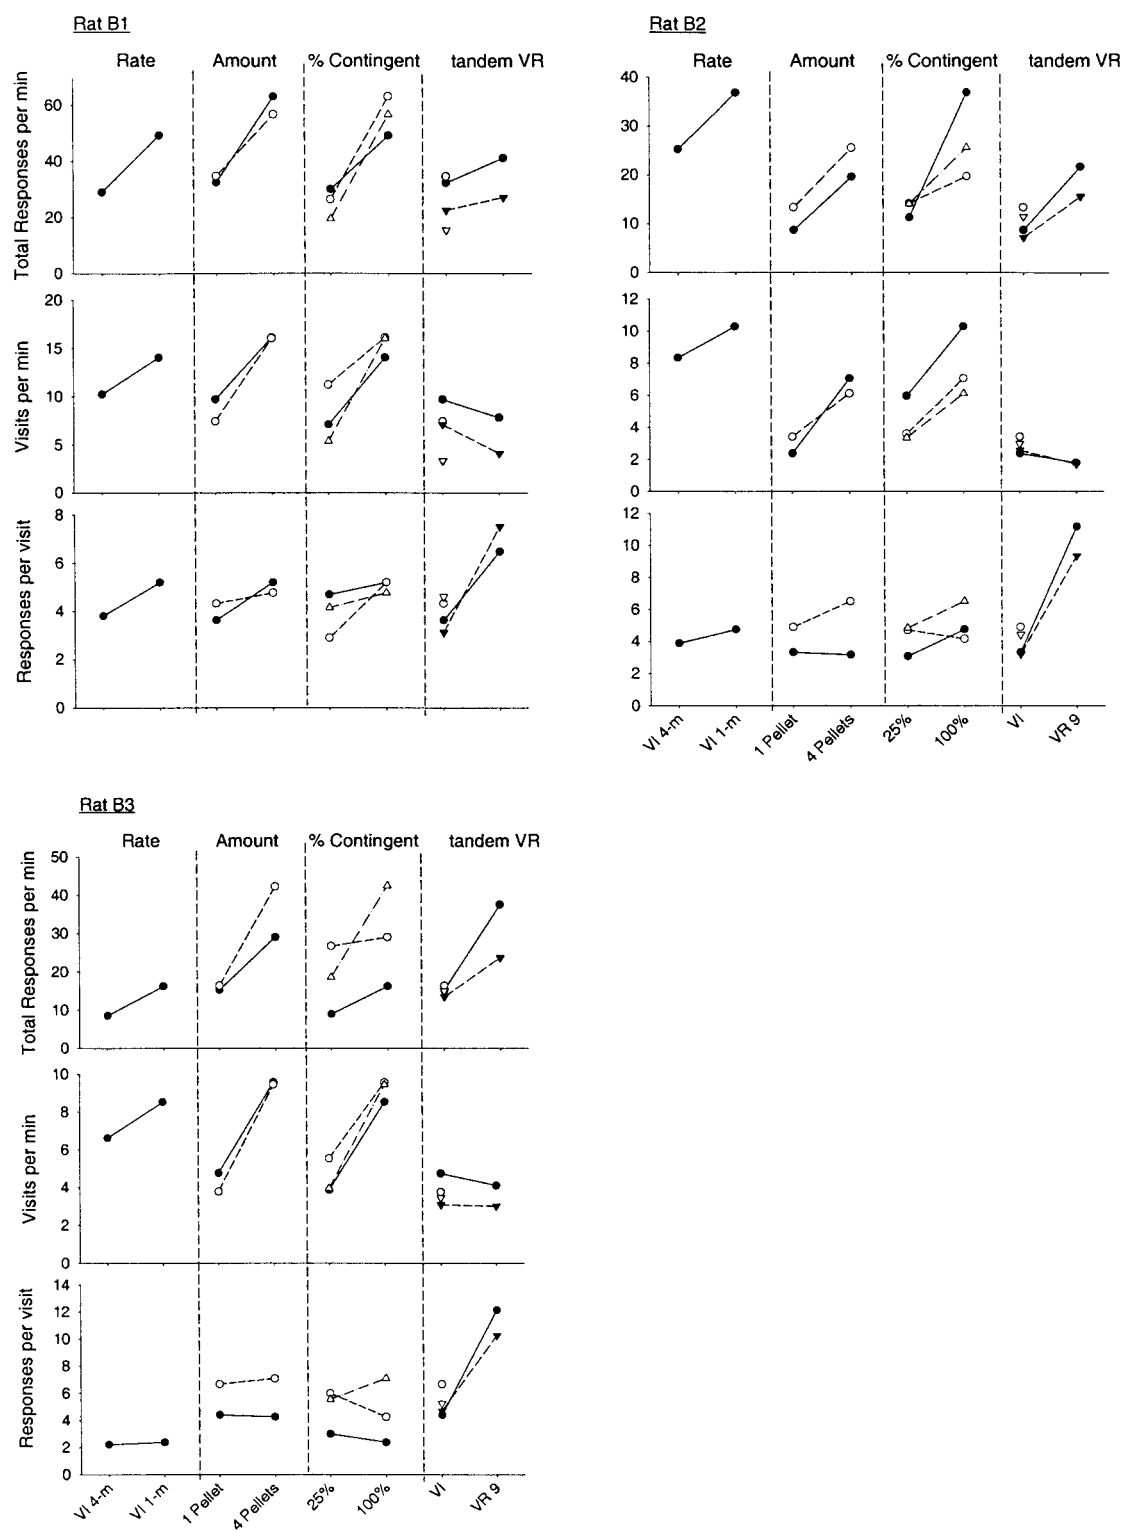


Fig. 8. The effects of different independent variables on total response rate, visit-initiation rate, and responses per visit for each of the 3 rats in Squad B. The plots are as described for Figure 7.




Fig. 9. Log survivor plots for the rats in Squad A, showing the persisting effects of the tandem VR 9 requirement on average responses per visit. The most important comparison is between the last session with the tandem VR requirement (dotted line) and the subsequent first session without the tandem VR 9 requirement (dashed line), and then between those two plots and the other two, which show the effects of extended exposure to no tandem VR 9 requirement. The plots are from the sessions indicated [as identified in Table 2: last session of Condition 3-Lean (solid); last session of Condition 5-Lean (dotted); first session of Condition 6-Lean (dashed); last session of Condition 6-Lean (dash, double dots)]. Proportions are plotted for consecutive 0.2-s intervals from 0 through 1 s; they are plotted for consecutive 1-s intervals from 1 s through 20 s.

Table 3

Average number of responses per visit (estimated from limbs of log survivor plots) for the sessions indicated. The values in the top and bottom rows show for each rat (Squads A and B) the effects of extended training under a VI 2-min schedule with no tandem VR requirement. The values in the row second from the top show the effects of training under a tandem VI 2 VR 9 schedule. The values in the row third from the top show the persisting effects of the tandem VR 9 during the first session without the VR 9.

Source of data (condition)	Rat						
	A1	A2	A3	A4	B1	B2	B3
Last session, no VR (C3-Lean)	2.37	2.86	1.53	4.86	3.49	3.01	4.45
Last session, VR 9 (C5-Lean)	6.10	7.14	12.06	11.79	6.48	12.45	11.66
First session, no VR (C6-Lean)	5.85	6.00	8.22	12.90	5.66	9.71	12.64
Last session, no VR (C6-Lean)	3.19	3.10	3.30	4.26	5.27	4.92	6.68

quirement increased the average number of responses per visit (i.e., visit length). Figure 10 shows a similar effect with a smaller tandem response requirement—either a tandem FR 4 (Condition 4-Lean) or a tandem VR 4 (Condition 8-Lean). Moreover, 11 of the 14 plots in Figure 10 show that visit length increased as a function of the size of the tandem response requirement over the range examined. Apparently, visit length is somewhat sensitive to the size of the response requirement, not merely to its presence or absence.

### DISCUSSION

A primary question is whether the log survivor analysis helps validate the two-mode conception of intermittently reinforced responding. The broken-stick appearance of the log survivor plots provides one bit of support. The two limbs plausibly result from two different processes or states, each correlated with a different rate of response emission. Indeed, in studies of naturally occurring behavior outside the laboratory, such broken-stick-appearing log survivor plots have often been taken as evidence that the behavior under investigation is organized into bouts (Clifton, 1987; Slater & Lester, 1982; Tolkamp & Kyriazakis, 1999). In addition, variables that had similar effects on total response rate had different effects on the slope and the intercept of the right limb of the log survivor plots (Figures 5 through 8). Apparently, those measures specify functionally distinct properties of responding that separately contribute to changes in total response rate. Our results thus support previous suggestions, cited in the introduction, that intermittently reinforced responding is usefully viewed as occurring in engagement bouts.

The method used here to estimate the rate and length of visits does not involve classifying individual responses as either visit-initiation or within-visit responses. It is therefore limited in the degree of detail that it can provide about the organization of intermittently reinforced responding. It does, however, indicate the rate and average length of visits. Such summary measures may be sufficient for addressing certain kinds of questions about behavior. That the effects were reasonably consistent among the individual rats and across replications increases our confidence

in the reliability and validity of the estimation technique. (See the Appendix for some additional considerations and data that bear on the validity of the estimates.)

A different approach to estimating the average number of responses per visit was described by Nevin and Baum (1980) and Baum (1992). They developed an equation, based on a plausible version of a two-mode conception of responding, that described how rate of reinforcement on a VI schedule varies as a function of rate of responding (i.e., the feedback function for VI schedules). They interpreted one of the free parameters of the equation as indicating the average number of responses per visit. Thus, by fitting the equation to a set of response- and reinforcement-rate data from a VI schedule, in which response rate is varied over a wide range (say, by satiation), a numerical estimate is obtained for the average number of responses per visit. It could be quite interesting to see how well the estimates based on the fits of Nevin and Baum's equation correspond to estimates from the limb of log survivor plots. Doing so, however, would require a different kind of data set from that generated in our study.

We do not know to what extent the patterns of responding that we have observed with key poking by rats generalize to other response forms (e.g., lever pressing by rats or key pecking by pigeons). Our hunch is that a more effortful response, such as pressing a fairly heavy lever, would generate longer and more variable within-visit IRTs. If so, the log survivor plots could show a more gradual bending instead of the kind of sharp break that was apparent in most of the log survivor plots generated by our rats. Such plots would complicate the estimation of visit-initiation rate and responses per visit. The relatively simple estimation procedure that we have described requires a reasonably linear right limb. Our hunch also is that relatively rich VI schedules might generate survivor functions whose limbs deviate more clearly and systematically from linearity than those generated by our rats. The delivery of food disrupts ongoing operant behavior and requires time for consumption. If these events occur frequently, the IRT distribution can be materially affected. There is a hint in the data (see Figure 5) that the right limbs are more nearly linear under the leaner (VI 4 min, open circles)



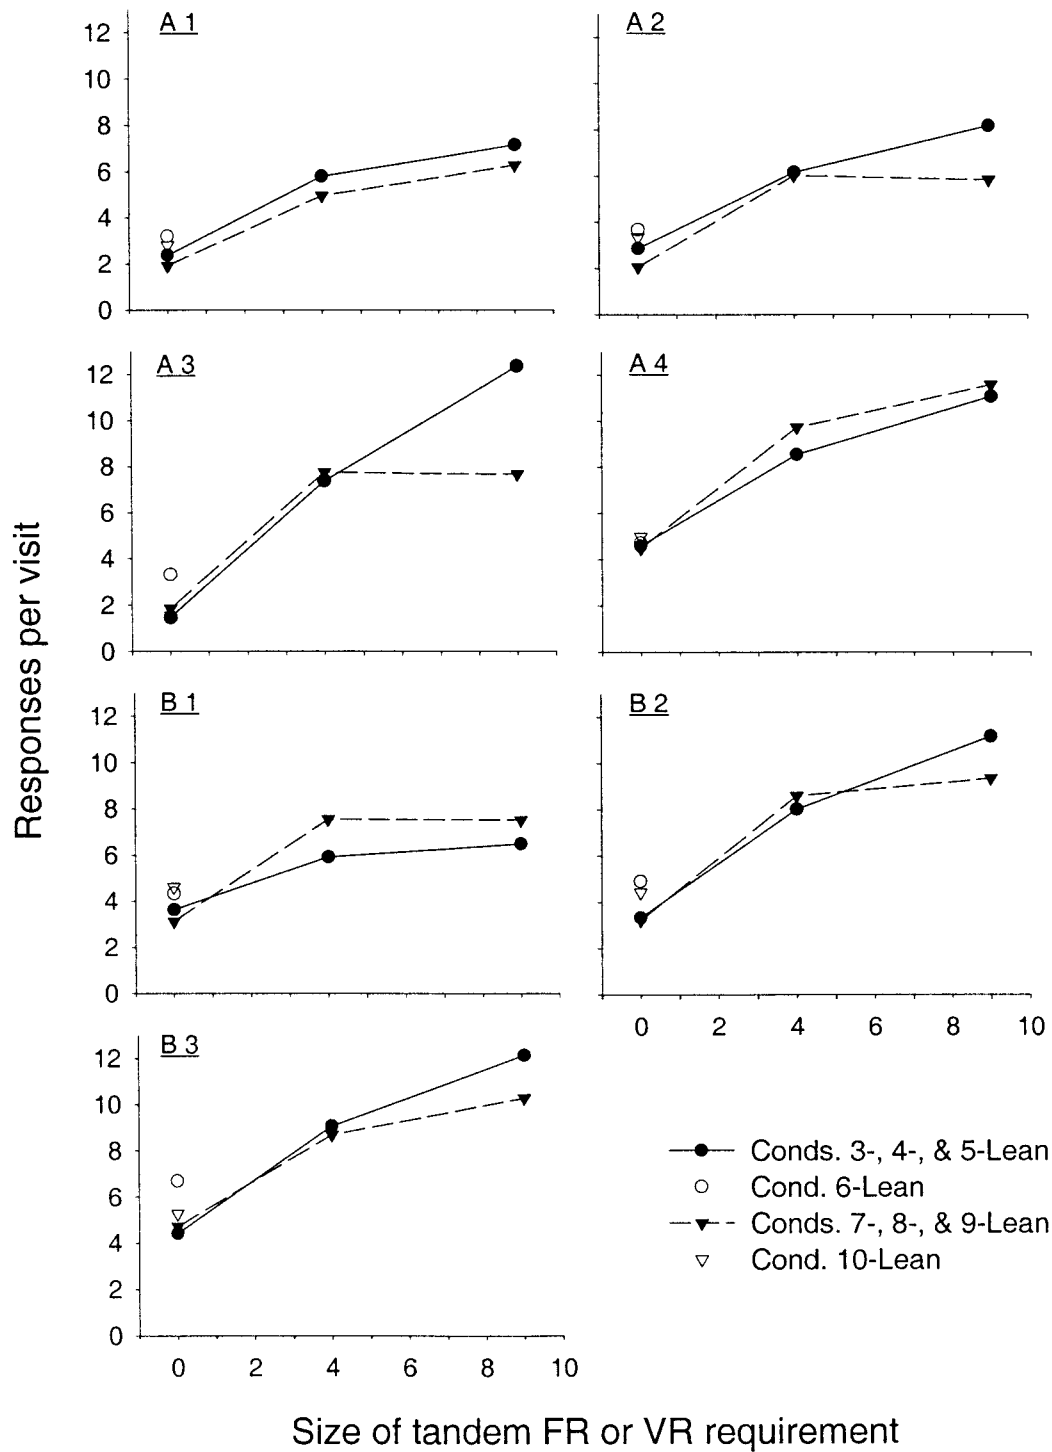


Fig. 10. Average number of responses per visit plotted as a function of the size of the tandem response requirement added to a VI 2-min schedule. The average numbers of responses per visit were estimated from the y-axis intercepts of the lines fitted to the limbs of the log survivor plots, as described in the text. Each point is the median value over the last three sessions of the indicated condition. For one series (filled circles), the four-response requirement was a tandem FR 4; for the other series (filled triangles), it was a VR 4. The legend indicates the conditions that generated the various points (see Tables 1 and 2).

than under the richer (VI 1-min, filled circles) schedules. Also, the simple version of the two-mode conception expressed in Figure 2 ignores possible sequential dependencies in responding, periodicities, additional kinds of disengagement states (see the Appendix), and other complexities. Figure 2 expresses a model of responding that surely is simplistic. Yet the method of estimating visit-initiation rate and average number of responses per visit that was based on this simple, first-approximation model was serviceable in revealing orderly, replicable results that differentiated classes of experimental variables in a way that the traditional unitary measure of response rate does not.

From the perspective of Herrnstein's (1970) hyperbolic equation for response rate, rate of reinforcement, amount of reinforcement, and percentage of response-contingent reinforcement are all variables that affect the relative reinforcement correlated with the designated response. It is of interest, then, that these three variables affected the total response rate of our rats in much the same way—namely, by affecting the slope of the log survivor limb (i.e., what we interpret as visit-initiation rate). The present results, then, could be interpreted as another illustration that relative reinforcement affects the disposition to engage in operant behavior (or the strength of operant behavior).

Adding a VR requirement in tandem to a VI schedule also increases the total response rate. It makes little sense, however, to think of the additional VR as increasing the relative reinforcement for the response. If anything, it might be thought to *decrease* the relative reinforcement because of the additional work or lengthened reinforcer delay timed from initiating a visit. We find it significant, therefore, that adding the tandem VR increased response rate by increasing the length of the visits despite often decreasing the rate of visit initiations. Apparently, not all increases in total response rate imply corresponding increases in relative reinforcement. And at least in this one case, the analysis of log survivor plots can distinguish the two kinds of effects.

The effect of the tandem VR requirement on responses per visit raises the possibility of a general approach toward understanding how different kinds of reinforcement schedules affect behavior. The basic idea is that

however else their effects might differ, different kinds of reinforcement schedules differ in the extent to which they favor the development of long visits. Under a VI schedule, for example, during a period of disengagement, the likelihood of reinforcement continually increases until a visit initiation occurs (cf., Anger, 1956; Cole, 1999; Morse, 1966). At that point, the likelihood that the next response will be reinforced immediately drops. Put casually, once the animal checks the availability of reinforcement by visiting, there is little point in lingering because the reinforcer is unlikely to have become available during the short interval since the last check. Thus, a response could become a discriminative event, signaling a period of reduced likelihood of reinforcer availability. The effect would be to generate short visits so that a relatively high proportion of responses on VI schedules would be visit initiations, making total response rate fairly sensitive to relative reinforcement variables (de Villiers & Herrnstein, 1976; Herrnstein, 1970; Williams, 1988).

In contrast, under an FR schedule each response produces an immediate increase in the likelihood of reinforcement (or an increase in the proximity to the next reinforcer). Thus, once a visit is initiated, the tendency should be to continue the visit until the reinforcer is obtained (Mazur, 1982, 1983). Upon delivery of the reinforcer, the likelihood that the next response will be reinforced drops immediately (i.e., the proximity to the next reinforcer immediately increases). Hence, the delivery of the reinforcer should be a discriminative event favoring disengagement from a visit. The well-known pattern of responding under FR schedules—a pause after reinforcement followed by a mostly uninterrupted period of responding until the next reinforcement—is consistent with this interpretation. Moreover, this interpretation anticipates that the average duration of the pause increases as a function of the size of the FR (Felton & Lyon, 1966; Mazur, 1982, 1983; Shull, 1979). The fixed response requirement imposes a delay between the visit initiation and reinforcement—a kind of changeover delay (Shull & Pliskoff, 1967). The larger the FR, the longer this delay will be. Because the rate of initiating a behavioral sequence decreases as a function of the delay

to reinforcement following the initiating response (cf. Fantino, 1977), the rate of visit initiations should decrease as a function of FR size (Shull, 1979). The average postreinforcement pause is, of course, simply the inverse of the rate (or speed) of the initial visit initiation following reinforcement; hence, the pause should increase as a function of FR size. Significantly, the log survivor plots of postreinforcement pauses from FR schedules have a long and reasonably linear right limb segment, the slope of which decreases toward zero as a function of FR size (see Capehart, Eckerman, Guilkey, & Shull, 1980; Shull, 1979, 1991). Also, the pause, viewed as the inverse of the rate of the initial visit initiation, would be expected to be sensitive to other relative reinforcement variables such as amount of reinforcement. That seems also to be the case (see Mazur, 1982; Shull, 1979).

Under VR schedules, the likelihood of reinforcement does not drop with each response as it does under VI schedules. Instead, the likelihood either remains the same (e.g., under random-ratio schedules) or increases (e.g., under arithmetic distributions). Thus, a response should not become discriminative for a period of reduced likelihood of reinforcement, and a visit, once engaged, should tend to continue until reinforcement. Because there are some small response counts (hence brief delays to reinforcement), the rate of visit initiation following disengagement should be fairly high (Mazur, 1983; Schlinger, Blakely, & Kaczor, 1990). Thus, performance under VR schedules should consist mainly of long visits between reinforcements and prompt reinitiations of visits whenever a disengagement occurs. Total response rate, being dominated by within-visit responding, should be insensitive to relative-reinforcement variables.

This approach to schedule effects may clarify why visit-initiation rate and the average number of responses per visit do not necessarily covary. If visit-initiation rate reflects the relative reinforcement correlated with the designated response, one might suppose that the disposition to remain engaged in a visit would do so as well. This reasoning led Nevin and Baum (1980) to assume that visit-initiation rate and visit length were highly positively correlated. Although our data showed evidence of a positive correlation under some

conditions (e.g., both tended to increase as a function of relative reinforcement variables), our data also showed that the average number of responses per visit changed far less in magnitude and consistency than visit-initiation rate. Moreover, under other conditions the two measures changed in opposite directions. Thus, on the surface our data could appear to contradict Nevin and Baum's assumption of a high positive correlation between those measures. We think, however, that their assumption is correct in its basic reasoning. The problem is that VI schedules, as discussed above, arrange contingencies that favor the development of brief visits, which would tend to reduce the correlation between the rate and length of visits. In this light, the modest tendency for the average number of responses per visit to increase as a function of relative reinforcement seems to provide striking support for the reasoning underlying Nevin and Baum's assumption even if it does not for their particular proposal about the degree of correlation. (By analogy, our argument here is similar to the suggestion that the rate of switching into a component of a concurrent VI VI schedule and the stay time in a component can be separately manipulated; e.g., MacDonall, 1998.)

There is an additional possible effect that we have ignored to this point. It is certainly possible that the kinds of variables that we have been considering also affect the within-visit response rate. If so, that would represent an additional contribution to changing total response rate. In principle, the log survivor analyses could be used to estimate within-visit response rates. One would fit a line to the initial left limb after subtracting the visit-initiation responses estimated by the fits to the right limb.<sup>3</sup> Although we tried this for some conditions and found (not surprisingly) somewhat higher within-visit response rates when the tandem VR 9 was in effect than when it was not, we have not pursued this analysis. The reason is that the very short IRTs are constrained by the resolving power of our apparatus and computer program.

We should note, however, that differences in within-visit response rate might turn out

<sup>3</sup> We thank Robert Kessel for alerting us to a similar approach for assessing the components of radioactive emission (Howard, 1963, pp. 14–15).

not to contribute much to differences in performance between different values of the same type of schedule and even between some different types of reinforcement schedules. The tempo of responding during periods of engagement assessed in different ways tends to be fairly insensitive to various manipulations (Davison & Charman, 1986; Mazur, 1982; Pear, Rector, & Legris, 1982; Rand, 1977). And when it does vary systematically, the effect is often traceable to differences in the proportion of responses that are unrecorded because of being, for example, insufficiently forceful (Elsmore & Brownstein, 1968) or inadequately located (Palya, 1992; Pear *et al.*, 1982).

The temporal pattern of responding can, of course, be modified by arranging time-based contingencies. For example, under a differential-reinforcement-of-low-rate (DRL) schedule, a response is reinforced only if it is preceded by a pause longer than some specified time interval (e.g., 5 s). Responding usually adjusts to such a contingency, in that pauses of approximately the required duration increase in frequency (see Lattal, 1991). Even though such responding does not appear to fit the kind of pause-burst pattern that has been the focus of this paper, an interpretation in terms of visits and within-visit engagement might, nonetheless, be appropriate for DRL performance. Several studies have arranged a procedure that is designed to separate the initiation of a timed pause from the timed pause itself. A rat, for example, presses one lever (Lever A) to start a DRL interval. Then, if the pause has been sufficiently long, a press on a second lever (Lever B) is reinforced. It turns out that the B-to-A times are quite sensitive to various reinforcement and motivation-like operations (e.g., deprivation level) but that the A-to-B times are rather insensitive to such operations (Mechner & Guevrekian, 1962; Migler & Brady, 1964). In other words, the tendency to initiate (i.e., visit) the DRL segment is more sensitive to such operations than is the learned pattern of responding during the visit. (See also Mechner, 1992, and Nevin, 1992, for analogous results, and see Galbicka & Platt, 1986, for a different approach toward an analogous conclusion.)

Perhaps the most important implication of the present results is that they clarify the con-

ditions under which total response rate—calculated in the usual way (i.e., total responses divided by total time)—will and will not vary sensitively as a function of relative reinforcement operations. To the extent that responding is dominated by visit initiations, total response rate should be sensitive to such operations. As discussed earlier, a higher proportion of responses should be visit initiations under VI schedules than under comparable VR schedules (and such indeed was the case in our data from VI and tandem VI VR schedules). Thus, total response rate on VI schedules should be a reasonable proxy of visit-initiation rate, yielding orderly functions of considerably generality, as has been well documented (de Villiers & Herrnstein, 1976; Herrnstein, 1970; Williams, 1988). It is possible that an even higher proportion of responses on VI schedules would be visit initiations if the operandum (e.g., a lever) were made relatively difficult to operate. Such preparations might yield total response rates that vary even more sensitively as a function of relative reinforcement operations. In contrast, response rates on VR schedules, particularly with an operandum that can be operated easily, should be relatively insensitive to relative reinforcement operations because such responding is dominated by (and thus the time saturated by) within-visit responding.

It has been known for some time that long IRTs are much more sensitive than short IRTs to various experimental operations that could be construed as affecting response strength (Blough, 1963; Schaub, 1967; Shull & Brownstein, 1970). From the present perspective, responses following relatively long IRTs will tend to be visit initiations, and those following the shortest IRTs will tend to be within-visit responses. Thus, the present results are entirely consistent with (and perhaps clarify the basis of) those earlier observations regarding the differential sensitivity of long versus short IRTs. Analogously, changes in the speed of a rat's running in an alley as a function of reinforcement and motivational variables turn out to result mainly from changes in the duration of breaks between bouts of running. There is little systematic change in the running speed once engaged (Cotton, 1953; Drew, 1939). Again, we interpret such results as indicating an effect of reinforce-

ment and motivational variables on the tendency to initiate visits (here, to the running mode).

## REFERENCES

- Anger, D. (1956). The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 52, 145–161.
- Ator, N. A. (1991). Subjects and instrumentation. In I. H. Iversen & K. A. Lattal (Eds.), *Techniques in the behavioral and neural sciences: Vol. 6. Experimental analysis of behavior* (Part 1, pp. 1–62). Amsterdam: Elsevier.
- Baum, W. M. (1992). In search of the feedback function for variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 57, 365–375.
- Berday, M. (1993). Defining bouts of behaviour: A three-process model. *Animal Behaviour*, 46, 387–396.
- Blough, D. S. (1963). Interresponse time as a function of continuous variables: A new method and some data. *Journal of the Experimental Analysis of Behavior*, 6, 237–246.
- Capehart, G. W., Eckerman, D. A., Guilkey, M., & Shull, R. L. (1980). A comparison of ratio and interval reinforcement schedules with comparable interreinforcement times. *Journal of the Experimental Analysis of Behavior*, 34, 61–76.
- Clifton, P. G. (1987). Analysis of feeding and drinking patterns. In F. M. Toates & N. E. Rowland (Eds.), *Feeding and drinking* (pp. 19–35). Amsterdam: Elsevier.
- Cole, M. R. (1999). Molar and molecular control in variable-interval and variable-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 71, 319–328.
- Cotton, J. W. (1953). Running time as a function of amount of food deprivation. *Journal of Experimental Psychology*, 46, 188–198.
- Davison, M., & Charman, L. (1986). On the measurement of time allocation on multiple variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 46, 353–362.
- de Villiers, P., & Herrnstein, R. J. (1976). Toward a law of response strength. *Psychological Bulletin*, 83, 1131–1153.
- Drew, G. C. (1939). The speed of locomotion gradient and its relation to the goal gradient. *Journal of Comparative Psychology*, 27, 333–372.
- Elmsmore, T. F., & Brownstein, A. J. (1968). Effort and response rate. *Psychonomic Science*, 10, 313–314.
- Fantino, E. F. (1977). Conditioned reinforcement: Choice and information. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 288–339). Englewood Cliffs, NJ: Prentice Hall.
- Felton, M., & Lyon, D. O. (1966). The post-reinforcement pause. *Journal of the Experimental Analysis of Behavior*, 9, 131–134.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. New York: Appleton-Century-Crofts.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 5, 529–530.
- Galbicka, G., & Platt, J. R. (1986). Parametric manipulation of interresponse-time contingency independent of reinforcement rate. *Journal of Experimental Psychology: Animal Behavior Processes*, 12, 371–380.
- Gilbert, T. F. (1958). Fundamental dimensional properties of the operant. *Psychological Review*, 65, 272–282.
- Hantula, D. A. (1991). A simple BASIC program to generate values for variable-interval schedules of reinforcement. *Journal of Applied Behavior Analysis*, 24, 799–801.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13, 243–266.
- Howard, R. A. (1963). *Nuclear physics*. Belmont, CA: Wadsworth.
- Killeen, P. (1969). Reinforcement frequency and contingency as factors in fixed-ratio behavior. *Journal of the Experimental Analysis of Behavior*, 12, 391–395.
- Lattal, K. A. (1991). Scheduling positive reinforcers. In I. H. Iversen & K. A. Lattal (Eds.), *Techniques in the behavioral and neural sciences: Vol. 6. Experimental analysis of behavior* (Part 1, pp. 87–134). Amsterdam: Elsevier.
- Logan, F. A. (1960). *Incentive: How the conditions of reinforcement affect the performance of rats*. New Haven, CT: Yale University Press.
- Luce, R. D. (1986). *Response times: Their roles in inferring elementary mental organization*. New York: Oxford University Press.
- MacDonall, J. S. (1998). Run length, visit duration, and reinforcers per visit in concurrent performance. *Journal of the Experimental Analysis of Behavior*, 69, 275–293.
- Mazur, J. E. (1982). A molecular approach to ratio schedule performance. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Matching and maximizing accounts* (pp. 79–110). Cambridge, MA: Ballinger.
- Mazur, J. E. (1983). Steady-state performance on fixed-, mixed-, and random-ratio schedules. *Journal of the Experimental Analysis of Behavior*, 39, 293–307.
- McGill, W. J. (1963). Stochastic latency mechanisms. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 309–360). New York: Wiley.
- Mechner, F. (1992). *The revealed operant: A way to study the characteristics of individual occurrences of operant responses*. Cambridge, MA: Cambridge Center for Behavioral Studies.
- Mechner, F., & Guevrekian, L. (1962). Effects of deprivation upon counting and timing in rats. *Journal of the Experimental Analysis of Behavior*, 5, 463–466.
- Mellgren, R. L., & Elmsmore, T. F. (1991). Extinction of operant behavior: An analysis based of foraging considerations. *Animal Learning & Behavior*, 19, 317–325.
- Migler, B., & Brady, J. V. (1964). Timing behavior and conditioned fear. *Journal of the Experimental Analysis of Behavior*, 7, 247–251.
- Morse, W. H. (1966). Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 52–108). New York: Appleton-Century-Crofts.
- Nevin, J. A. (1992). The discriminated operant in pigeon multiple schedule performance: Effects of reinforcement rate and extinction. In F. Mechner (Ed.), *The revealed operant: A way to study the characteristics of individual occurrences of operant responses* (pp. 71–76). Cambridge, MA: Cambridge Center for Behavioral Studies.
- Nevin, J. A., & Baum, W. M. (1980). Feedback functions



- for variable-interval reinforcement. *Journal of the Experimental Analysis of Behavior*, 34, 207–217.
- Nevin, J. A., & Grace, R. C. (2000). Behavioral momentum. *Behavioral and Brain Sciences*, 23, 73–130. (includes commentary)
- Palya, W. L. (1992). Dynamics in the fine structure of schedule-controlled behavior. *Journal of the Experimental Analysis of Behavior*, 57, 267–287.
- Pear, J. J., & Rector, B. L. (1979). Constituents of response rate. *Journal of the Experimental Analysis of Behavior*, 32, 341–362.
- Pear, J. J., Rector, B. L., & Legris, J. A. (1982). Toward analyzing the continuity of behavior. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Matching and maximizing accounts* (pp. 3–24). Cambridge, MA: Ballinger.
- Peele, D. B., Casey, J., & Silberberg, A. (1984). Primacy of interresponse-time reinforcement in accounting for rate differences under variable-ratio and variable-interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 149–167.
- Premack, D. (1965). Reinforcement theory. In D. Levine (Ed.), *Nebraska symposium on motivation* (pp. 123–180). Lincoln: University of Nebraska Press.
- Rand, J. F. (1977). Behaviors observed during S– in a simple discrimination learning task. *Journal of the Experimental Analysis of Behavior*, 27, 103–117.
- Schaub, R. E. (1967). Analysis of interresponse times with small class intervals. *Psychological Record*, 17, 81–89.
- Schlinger, H., Blakely, E., & Kaczor, T. (1990). Pausing under variable-ratio schedules: Interaction of reinforcer magnitude, variable-ratio size, and lowest ratio. *Journal of the Experimental Analysis of Behavior*, 53, 133–139.
- Shull, R. L. (1979). The postreinforcement pause: Some implications for the correlational law of effect. In P. Harzem & M. D. Zeiler (Eds.), *Advances in analysis of behaviour: Vol. 1. Reinforcement and the organization of behaviour* (pp. 193–221). Chichester, England: Wiley.
- Shull, R. L. (1991). Mathematical description of operant behavior: An introduction. In I. H. Iversen & K. A. Lattal (Eds.), *Techniques in the behavioral and neural sciences: Vol. 6. Experimental analysis of behavior* (Part 2, pp. 243–282). Amsterdam: Elsevier.
- Shull, R. L., & Brownstein, A. J. (1970). Interresponse time duration in fixed-interval schedules of reinforcement: Control by ordinal position and time since reinforcement. *Journal of the Experimental Analysis of Behavior*, 14, 49–53.
- Shull, R. L., & Pliskoff, S. S. (1967). Changeover delay and concurrent schedules: Some effects on relative performance measures. *Journal of the Experimental Analysis of Behavior*, 10, 517–527.
- Sibley, R. M., Nott, H. M. R., & Fletcher, D. J. (1990). Splitting behaviour into bouts. *Animal Behaviour*, 39, 63–69.
- Slater, P. J., & Lester, N. P. (1982). Minimising errors in splitting behaviour into bouts. *Behaviour*, 79, 153–161.
- Tolkamp, B. J., Allcroft, D. J., Austin, E. J., Nielsen, B. L., & Kyriazakis, I. (1998). Satiety splits feeding behaviour into bouts. *Journal of Theoretical Biology*, 194, 235–250.
- Tolkamp, B. J., & Kyriazakis, I. (1999). To split behaviour into bouts, log-transform the intervals. *Animal Behaviour*, 57, 807–817.
- Walter, D. E., & Palya, W. L. (1984). An inexpensive experiment controller for stand-alone applications or distributed processing networks. *Behavior Research Methods, Instruments, & Computers*, 16, 125–134.
- Williams, B. A. (1988). Reinforcement, choice, and response strength. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology* (2nd ed., Vol. 2, pp. 167–244). New York: Wiley.

Received August 14, 2000  
Final acceptance February 22, 2001

## APPENDIX

### VALIDITY OF THE ESTIMATES OF VISIT-INITIATION RATE AND THE AVERAGE NUMBER OF RESPONSES PER VISIT

With the simulations (Part 1) it was possible to check the validity of the estimates that were derived from the fits to the limbs of log survivor functions. The simulation program coded the responses to indicate whether they were generated by the visit-initiation routine or by the within-visit routine. Thus, we could compare the estimates that were derived from the log survivor fits with the values that were computed directly from the coded response totals (Figure 4).

With the data from the rats, in contrast, we have no way to assess visit-initiation rate and average number of responses per visit independently of the fits to the limbs of the log survivor plots. Thus, we cannot determine for the rats' data, as we could with the simulations, how well the estimates correspond to actual values. Indeed, with the rats' data the classification of visit-initiation versus within-visit responses is a theoretical construction.

The following lines of evidence have helped to convince us that this classification and our estimation methods are useful: (a) The log survivor plots were nearly always approximations of a broken stick, suggesting two processes; (b) the downward extent of the left limb and the slope of the right limb were visibly altered by different classes of independent variables; (c) numerical estimates of those characteristics of the log survivor plots (i.e., the slope and y-axis intercept that were derived from fitting a line to a segment of the right limb) likewise varied systematically as a function of the different classes of independent variables; (d) these estimates were reasonably consistent over replications

and among the rats; and (e) the grouping of independent variables on the basis of their effects on the estimates corresponded well to the grouping of those variables based on other influential theoretical accounts (e.g., Herrnstein, 1970) and on common sense.

An additional approach toward validating the estimation procedure is to see how well total response rate, computed from session totals in the usual way (i.e., total number of responses divided by time), can be predicted from combining the estimates of visit-initiation rate and average number of responses per visit. If the two-mode conception, as diagrammed in Figure 2, provides an accurate description, then total response rate is simply the product of visit-initiation rate and average number of responses per visit. We can assess the accuracy of the two-mode conception, then, by seeing how well the product of our two estimates correlates with the measure of total response rate that is computed from session totals. Because of time taken during the visits, the estimates of visit rate would need to be adjusted downward, as described earlier in relation to Figure 4, for the predictions to hold.

For the adjustment, we needed an estimate of the average IRT duration for within-visit responding for each rat. We used the median IRT from the last session of Condition 5, the condition that arranged the tandem VR 9 schedule. This condition generated log survivor plots whose initial limbs dropped well below a proportion of .5 (i.e., the median) before the break to the right limb. Thus, we could be confident that the median IRT from this condition represented mostly within-visit IRTs. We do not know how well this median IRT corresponded to the medians from other conditions. If there were large differences, the adjustment ideally should incorporate the IRTs from each condition separately. As a first approximation, however, and for the sake of simplicity, we used the single median IRT for each rat. The point of the adjustment was to reduce the visit-initiation rate, as estimated by the slope of the limb segment, in accordance with the fact that only a portion of the time was available for visit initiations because of time occupied by the visits. The average time spent per visit is estimated by  $NW$ , where  $N$  is the average number of responses per visit and  $W$  is the average within-visit IRT (here, the

median from the last session of Condition 5). The time available for a visit is the average disengagement time, or  $1/V$ , where  $V$  is the visit-initiation rate as estimated by the slope of the limb. The proportion of the session time available for visit initiations, then, is  $(1/V)/[(1/V) + NW]$ . Thus, the predicted total response rate,  $R'$ , from the log survivor estimates is

$$R' = NV \left( \frac{(1/V)}{(1/V) + NW} \right) \\ = \frac{N}{(1/V) + NW}. \quad (2)$$

With some substitution and rearranging of terms, this formula is equivalent to Equations 5 and 6 in Shull (1991), which also describe the relation between observed response rate and visit-initiation rate under the assumption of a two-mode conception. Those equations were generated from somewhat different considerations from those that led to Equation 2 here.

Sometimes the product of the adjusted estimates corresponded closely to response rate computed from session totals (e.g., top panel in Figure 11), but sometimes it did not (lower panel of Figure 11). Table 4 gives values for all 7 rats of the slope, intercept, and  $r^2$  of the best fitting straight line through scatter plots like those shown in Figure 11. Also shown for each rat is the IRT duration that was used to adjust the visit-initiation rate. The  $r^2$  values indicated good fits for most of the rats, and the slopes for several were moderately close to 1.0. For those rats, the two-mode conception (as diagrammed in Figure 2) appears to provide a reasonable, first-approximation description of responding. For some other rats, however (especially Rats A3 and A4), the slopes were substantially higher than 1.0, and the  $r^2$  values indicated poorer fits. For those rats, the two-mode conception (as diagrammed in Figure 2) does not provide an adequate description of responding. Moreover, the slopes for all rats were above 1.0, suggesting systematic effects that are not represented in the two-mode model or in our estimation methods, even for the rats whose fits were otherwise fairly good.




Fig. 11. Total response rate estimated from the product of visit-initiation rate and average number of responses per visit is plotted as a function of response rate computed the usual way from session totals for 2 rats. For these estimates visit-initiation rate was adjusted downward for time spent engaged in visits, as described in the text. For each rat there were 6 points per condition (i.e., 3 sessions per condition  $\times$  2 discriminative stimuli) and 10 conditions. The dashed diagonal line indicates perfect correspondence between the two measures of total response rate (i.e., slope = 1.0). The data set in the upper panel (Rat B1) is labeled "best" in the sense that it had the second highest value of  $r^2$  and the closest slope to 1.0 among the 7 rats. The data set in the lower panel (Rat A4) is labeled "worst" in the sense that it generated the lowest value of  $r^2$  among the 7 rats. The best fitting linear equation (Pearson product moment) for each data set is shown in the corresponding panel, along with the  $r^2$  value and the median IRT used to adjust visit-initiation rate.

Table 4

The results of straight-line fits (Pearson product-moment) to scatter plots relating response rate computed from session totals ( $x$  axis) to the product of visit-initiation rate and average number of responses per visit estimated from log survivor plots (e.g., Figure 11). For each rat, there were 6 points per condition (i.e., 3 sessions per condition  $\times$  2 discriminative stimuli) and 10 conditions. Shown are the slope,  $y$ -axis intercept, and  $r^2$  values for each fit. For the fits in the left columns, the estimated visit-initiation rate was adjusted downward for time spent during the visit (adjustment set at the median IRT from the last session of Condition 5, listed in the fifth column); for the fits in the right columns, the estimated visit-initiation rate was not adjusted.

Rat	Visit-initiation rate adjusted			Median IRT (min)	Visit-initiation rate not adjusted		
	Slope	Intercept	$r^2$		Slope	Intercept	$r^2$
A1	1.25	0.39	.89	0.0042	1.54	-2.13	.89
A2	1.37	-5.71	.94	0.0052	2.30	-21.06	.89
A3	1.83	1.94	.84	0.0037	2.11	1.24	.84
A4	1.39	3.09	.66	0.0035	1.81	-0.22	.66
B1	1.09	-1.23	.96	0.0033	1.49	-9.17	.96
B2	1.20	-0.83	.92	0.0037	1.48	-3.61	.92
B3	1.17	1.46	.97	0.0037	1.63	-4.3	.97

Table 4 also lists the slope, intercept, and  $r^2$  values for corresponding scatter plots but for which the visit-initiation rates were not adjusted downward. There was no apparent effect on the values of  $r^2$  as a result of adjusting or not adjusting visit-initiation rate. But the slopes were higher in every case when the initiation rate was not adjusted, which simply confirms the earlier point that the unadjusted visit-initiation rate will overestimate the visit-initiation rate that actually occurs and, thus, will overestimate the total response rate.

There are a number of factors that could have caused the systematic overestimation of response rate and unsystematic variability (i.e., slopes above 1.0 and relatively low values of  $r^2$ ) such as occurred especially for Rats A3 and A4. These include (a) systematic deviations from linearity of the limbs of the log survivor plots, (b) insufficient adjustment for time spent during a visit, and (c) the occurrence of additional response modes beyond the two represented in Figure 2 (e.g., Berdoy, 1993; Tolkamp, Allcroft, Austin, Nielsen, & Kyriazakis, 1998).

This possibility seems especially likely to be the case for Rat A4, whose scatter plot is shown in the lower panel of Figure 11, and for Rat A3. For these rats the limbs of the log survivor plots appeared to be reasonably linear (see, e.g., the samples in Figure 5). Increasing the duration of the IRT used to adjust visit-initiation rate much beyond the one that we actually used caused the scatter plots

to bend downward, thus further reducing the  $r^2$  values. There was, however, some evidence for these 2 rats that suggested a third behavioral mode of extended disengagement. During some conditions, a high proportion of cycles contained no nose poking at all despite an overall fairly high response rate. For example, during Condition 5 (last three sessions) 23% of the 50-s components for Rat A4 and 26% for Rat A3 contained no responses (compared to 1% and 5% for Rats A1 and A2, respectively). Such no-response components were not incorporated into the IRT distributions but did add to the time base for computing total response rate. If these no-response components represented a third responding mode, the product of our estimates of visit-initiation rate and responses per visit would predict higher values of total response rate than what actually occurred, an implication that is consistent with the slope (Table 4) being substantially greater than 1.0 for Rats A3 and A4. When we observed the rats, these 2 rats seemed to spend unusually large blocks of time sniffing at the rear of the chamber, appearing to be thoroughly disengaged from the reinforcement schedule.

Other third-mode-like periods may have occurred more generally and might have contributed to the slopes being higher than 1.0 for all rats. For example, the rats were often observed to take an unusually long time to resume responding following consumption of the multipellet reinforcers during the rele-

vant conditions. Perhaps factors of this sort are responsible for the high slope for Rat A2.

The evidence in Figure 11 and Table 4 shows that the two-mode model (Figure 2) is an incomplete description of our rats' responding (although we find the correspondence for several rats remarkable given the simplicity of the assumptions and the fact that the predictions were generated by fits to a small segment of the log survivor plots plus the median IRT from one session). Nonetheless, we again raise the possibility that the un-

adjusted visit-initiation rate might be a particularly useful measure of the disposition to engage in operant behavior because it is relatively unconstrained (Shull, 1991). In any case, we note that none of the data presented in this section compromise two points that we consider significant. The two performance measures that we interpret as estimates of visit-initiation rate and number of responses per visit varied systematically as a function of different classes of independent variables. Also, total response rate is a composite measure of different classes of performance.